

APPLICATION FOR LETTERS PATENT OF THE UNITED STATES

CERTIFICATE OF MAILING "EXPRESS MAIL"	
"Express Mail" Mailing Label Number	<u>EL 738966 070 US</u>
Date of Deposit	<u>April 12, 2001</u>
I hereby certify that this paper or fee is being deposited with the United States Postal Service "EXPRESS MAIL POST OFFICE TO ADDRESSEE" Service under 37 CFR 1.10 on the date indicated above and is addressed to the Assistant Commissioner for Patents, Washington, D.C. 20231.	
<u>SALLIE SPICER</u> (type or print name of person certifying)	
<u>Sallie Spicer</u> Signature	

SPECIFICATION

To all whom it may concern:

Be It Known, That I, WAYNE L. HUTCHINSON, of Columbia, SC, have invented certain new and useful improvements in METHODS FOR DYNAMICALLY CONFIGURING THE CARDINALITY OF KEYWORD ATTRIBUTES, of which I declare the following to be a full, clear and exact description:

METHODS FOR DYNAMICALLY CONFIGURING THE CARDINALITY OF KEYWORD ATTRIBUTES

Wayne L. Hutchinson

5 FIELD OF THE INVENTION

The present invention relates to methods for dynamically expanding the cardinality associated with keyword attributes in a data store.

BACKGROUND OF THE INVENTION

10 Information collected about customers in the retail industry is typically stored by individual retailers in one or more databases. The information is used by the retailer to facilitate communications with their customers, to assist with the billing of their customers, and to improve the marketing of products or services to their customers. The warehousing
15 and mining of this information is often referred to as customer relationship management (CRM). A plethora of information technology services and tools have been created in recent years to assist retailers in developing CRM systems from their existing customer databases.

Many retailers are still finding the task of creating CRM systems
20 daunting, because often each retailer will have multiple disparate legacy databases having information about customers, and these databases need to be integrated in order to achieve a complete customer data representation. Moreover, the quality and accuracy of any CRM system is directly dependent upon the completeness of the customer data on which it depends.
25 Accordingly, one of the paramount prerequisites of any successful CRM system is to establish a centralized data representation of a customer.

To achieve this, requires a single database table with design consistency such that all the various permutations of similar customer attributes are identified within the table as a single consistent customer attribute. For example, consider two databases within a retail establishment, the first is used to record and market to customers purchasing by mail order catalogue, and the second is used to record and market to customers purchasing by phone.

Over the course of time, each of these databases may have been populated, enhanced, and utilized by two separate business units within the retail establishment. As a result, each database table may represent attributes associated with the same customer in mutually exclusive manners making the integration of the two databases a labor intensive and expensive process for the retail establishment.

For instance, if the first database table recorded the various phone numbers of a customer as fields or attributes with field names such as "phone1," "phone2," "phone3," and the like, while the second database table recorded the same phone numbers of the customer as fields or attributes with field names such as "home," "cell," "work," and the like, then creating a centralized representation of the customer would require mapping the disparate naming conventions to a single consistent naming convention. Moreover, if each database has any significant number of attributes associated with each customer within the database, then the mapping of the disparate attributes of the two databases to a consistent schema may be very expensive for the retail establishment. Furthermore, achieving a consistent naming convention for disparate database tables is largely a manual process which typically cannot be achieved through automated software processes.

As one skilled in the art will readily appreciate, as the number of disparate databases and disparate naming conventions for attributes of each database increases within a retail establishment, the expense of migrating to a consistent centralized database will correspondingly increase as well.

5 Therefore, many retail establishments have been slow to migrate towards the CRM paradigm, and they have not benefitted, to any significant degree, from the CRM services and tools available in the industry.

Accordingly, methods are needed wherein a single database table definition would permit the rapid and easy migration of disparate database
10 tables into a single consistent table. To achieve this, the cardinality (e.g., the relationship between the attribute and the value of the attribute, such as by way of example only, one to one, one to many, and the like) associated with the database attributes must be dynamically modifiable. By
15 permitting the attributes of a table definition, associated with a database table, to be dynamically configured a single table definition may be defined wherein disparate attributes are seamlessly integrated.

In this way, retail establishments and any other organization, may rapidly and efficiently integrate all the information within the establishment's disparate databases by migrating the databases to a single
20 table definition using standard automated database extraction and database insertion operations. This permits establishments the ability to rapidly create powerful and centralized data warehouses without the need to eliminate or modify any established information collection processes.

Database designers cannot adequately predict all of the attributes
25 that will ever be of interest, correspondingly a flexible design approach to accommodate any unforeseen attributes which develop over the course of time is desirable. The adaptability of a new approach, which would permit

a single database design to integrate with multiple disparate customer database environments without changing the single design, is highly desirable.

5 SUMMARY OF THE INVENTION

Accordingly, an object of the invention is to provide methods for dynamically configuring the cardinality associated with keywords in a data store. A table definition, associated with a table of a data store, may be defined to include a keyword value field which when combined with
10 keyword and identification fields may be used to form a composite key associated with the table.

The cardinality associated with any given keyword may be dynamically altered as unique values are added to the keyword value field, without the need to alter the table definition of the data store. In this way,
15 disparate data stores having information relevant to a single keyword may be combined in a single table definition that permits the keyword cardinality to be expanded without redesigning the table definition to achieve the same. As a result, data design consistency may be achieved programmatically without the need for labor intensive redesigning or ad
20 hoc creation of table definitions which are typically required when information from disparate tables are integrated into a single table.

Additional objectives, advantages and novel features of the invention will be set forth in the description that follows and, in part, will become apparent to those skilled in the art upon examining or practicing the
25 invention. The objects and advantages of the invention may be realized and obtained by means of the instrumentalities and combinations particularly pointed out in the appended claims. To achieve the foregoing

and other objects and in accordance with the purpose of the present invention, methods for dynamically configuring a cardinality associated with keyword attributes are provided.

One aspect of the present invention provides a method for
 5 dynamically configuring a cardinality of keyword attributes having executable instructions, wherein a table is received. The table has a table definition including an identification field, a keyword field, and a keyword value field. First and second values are inserted into the table and associated with a first keyword. The values are part of the keyword value
 10 field while the first keyword is part of the keyword field. Further, a first identification which is part of the identification field is associated with the first keyword and the values.

Another aspect of the present invention provides a method for
 expanding a table definition having executable instructions where a table
 15 definition is received. The table definition includes an identification field, a keyword field, and a keyword value field. First and second values of the keyword value field are associated with a first keyword of the keyword field.

Furthermore, a method for expanding a keyword by permitting one or
 20 more keyword values to be associated with each keyword is provided having executable instructions for receiving a table having an identification field, a keyword field, and a keyword value field. Moreover, a first keyword is received and associated with the keyword field. The first keyword has first and second values where each value is associated with
 25 the keyword value field.

Still other aspects of the present invention will become apparent to those skilled in the art from the following description of an exemplary

embodiment, which is by way of illustration, one of the best modes contemplated for carrying out the invention. As will be realized, the invention is capable of other different and obvious aspects, all without departing from the invention. Accordingly, the drawings and descriptions
5 are illustrative in nature and not restrictive.

BRIEF DESCRIPTION OF THE DRAWINGS

The accompanying drawings, incorporated in and forming part of the specification, illustrate several aspects of the present invention and,
10 together with their descriptions, serve to explain the principles of the invention. In the drawings:

Fig. 1 depicts a flow diagram of a method for configuring a cardinality of keyword attributes;

Fig. 2 depicts a flow diagram of a method for expanding a table
15 definition;

Fig. 3 depicts a schematic diagram of table definitions; and

Fig. 4 depicts a flow diagram of an access method having two tables;
and

Fig. 5 depicts a flow diagram of a method for expanding a keyword.
20

DETAILED DESCRIPTION

The present invention provides methods for dynamically configuring the cardinality associated with keyword attributes of a data store. One embodiment of the present invention is implemented by providing a single
25 table definition within the NCR Teradata database and utilizing standard SQL database operations to extract information from disparate databases which is then inserted into the single table definition.

However, as one skilled in the art will readily appreciate alternative database technologies and database application programming interface (API) languages (now known or hereafter developed) may also be readily employed. Moreover, as one skilled in the art will readily appreciate no
 5 database is needed at all, since the present invention could be written as a stand alone application providing its own database operations and its own database API.

Information gathered by organizations is typically stored in databases. A database is a set of software applications which requires the
 10 information to be stored in a structured manner into one or more data structures referred to as database tables. Each table within the database includes columns (e.g. fields) and rows (e.g. tuples or records). Typically, each table has a single key or a composite key where access to a row of the table may be acquired when the appropriate key is provided.

15 A single key consists of a single unique field within the table, such that no value within the unique field recurs within the table. A composite key includes two or more fields within the table, such that no set of values within the fields, used for purposes of key construction, are identical within the table.

20 Furthermore, the fields of the table are defined when the table is initially created, this schema is often referred to as a table definition. If new fields are later added to the table, which were not initially defined, then the table definition must be modified to accommodate the newly added fields. As one skilled in the art will readily appreciate, modifying a table
 25 definition is an expensive and undesirable operation for large databases, requiring the table indexes and the entire database utilizing the table to be reconfigured or rebuilt. Also, during this operation access to the database

is often restricted and available computing cycles for other computing operations become scarce. As a result, large databases are typically rebuilt on dedicated computing devices or are rebuilt during off-peak usage hours, so a user's response time is not measurably affected.

5 Accordingly, great effort and time are usually expended in the designing and creation of a database table definition, so that the definition may last for a reasonable amount of time without the need for modification. Moreover, when multiple databases are migrated to a single database a single table definition will need to be created to achieve consistency among
10 all fields which are logically the same but stored under different field names within the multiple databases.

 Consider a table definition having three fields: Identification, Keyword, and Value; where the "Identification" field includes values associated with uniquely identifying a specific customer of a retail
15 establishment, these values may include, by way of example only, account numbers, social security numbers, last names, and the like. Moreover, an Identification field may be manufactured by additional software, such that a specific customer is resolved and given a unique Identification number by the additional software. This may be done, by way of example only, by
20 analyzing information received from an unknown customer, and determining based on the information, such as address, last name, and/or social security number, that the customer in question is uniquely identifiable with a Identification number "n." Of course, as one skilled in the art will appreciate any number of data mining techniques may be used
25 to identify a specific customer and assign a unique Identification number to that specific customer.

Furthermore the "Keyword" field includes string values which are the field names of other disparate database tables, such as by way of example only "phone," "address," "cell," "name," and others. Moreover, the "Value" field includes values associated with the keyword field. For example, if a keyword field value consists of the string "name" then a corresponding value field may include an associated value of "Hutchinson."

By permitting the fields of a disparate database to be included as string values in the keyword field, then a single table definition having three fields as presented above may be used to permit the integration of multiple disparate database tables into the single integrated table definition presented above. In this way, a retail establishment having separate or conflicting field names within disparate databases for information regarding the same customer may be seamlessly migrated to a single table using the table definition presented. Standard database extraction operations may then be used to extract the fields of a database associated with a customer and used in connection with standard database insertion operations to insert the extracted fields and their corresponding values into the table definition presented.

Fig. 1 depicts one flow diagram of a method for configuring a cardinality of keyword attributes. Initially, a table definition is received in step 10 having three fields defined: an "Identification" field defined in step 20, a "Keyword" field defined in step 30, and a "Value" field defined in step 40. The Identification field may include any information which is used to uniquely identify an account, customer, client, and the like. The Keyword field includes strings associated with fields of external data stores, and the Value field includes a value for a corresponding Keyword string. Although, as one skilled in the art will appreciate, no Identification

field is required at all since this field may also be listed as a string name within the Keyword field itself.

The three fields are associated with a table in step 50 by creating a table definition. The table is created having a composite key which includes all three fields of the table definition in step 60. An initial Keyword string ("keyword1") is acquired from one or more external data stores along with an initial Value ("value1"). As previously presented, the Keyword string (e.g., keyword1) is a field name of an external table definition and the Value associated with the string (e.g., "value1").

Acquisition of keyword1 and value1 may be accomplished, by way of example only, by extracting these items from an external database using any standard SQL search query.

In step 70, keyword1 and value1 are inserted into the table, so as to create a first tuple or row in the table in step 80. Furthermore, keyword1 is also associated with a second value ("value2") and inserted in the table in step 90, so as to create a second tuple or row in the table in step 100.

Assume for purposes of illustration only, a customer of a retailer with an account number of "11111111." However, as previously presented the account number may just be used by another piece of software to assign a unique number to a customer associated with the account number. Or, alternatively the account number itself may be used as the Identification number, as discussed above. Further, the retailer maintains two disparate databases on customers, these databases are not integrated. Each database uniquely identifies the customer with the above account number, and each database contains a field entitled "phone" for the customer, however, in the first database a value associated with the phone field contains the

customer's home phone number, while the second database includes a phone field's value containing the customer's work phone number.

In the present example, a table definition is established having three fields and two rows. The fields include "Identification," "Keyword," and "Value." Keyword1 is the string "phone," and it is associated with the Keyword field. Moreover, the Identification field includes the account number "11111111." The Value field has two entries which represent the customer's home phone number and work phone number. The resulting table has a first row including the tuple "11111111, phone, home-phone-number" and a second row including the tuple "11111111, phone, work-phone-number." Access to the table is obtained using a unique composite key which includes all three fields of the table. The table appears as follows:

<u>Identification</u>	<u>Keyword</u>	<u>Value</u>
11111111	phone	home-phone-number (e.g., 000-0000)
11111111	phone	work-phone-number (e.g., 111-2222)

In this way two disparate databases were integrated into a single database table definition where the Keyword phone which originally had a cardinality of one-to-one (e.g., the Keyword field value "phone" occurred only once within each database and had only one possible value, namely home or work phone number). Yet, by migrating to a table definition, as presented above, the cardinality associated with the keyword field is dynamically expanded to include a many-to-one relationship (e.g., the resulting table includes two phone Keywords, with each Keyword has one possible value, namely home or work phone number). Two disparate databases are merged into a single database table and the effective

cardinality originally associated with fields of the disparate databases are dynamically modified during the transition.

As one skilled in the art will readily appreciate, the method depicted in Fig. 1 will readily permit automatic extraction and migration of multiple disparate database tables into a single integrated database table, thereby permitting a single centralized data store wherein the fields may be modified so as to change a cardinality associated therewith. This permits field extensibility and better data mining, since all information regarding a customer, may be centralized for purposes of data indexing, searching, retrieving, and extracting.

Fig. 2 depicts one flow diagram of a method for expanding a table definition. A table definition used to define the attributes of a data store table is received in step 110. The table definition defines a single integrated table to include an "Identification" field in step 120, a "Keyword" field in step 130, and a "Keyword Value" field in step 140. A first value is identified in step 170 and a second value identified in step 190. The first and second values are not equal, and each value is independently associated with a first keyword in step 160.

The association between the first keyword and the first value creates a first table entry in step 150, and the association between the second keyword and the second value creates a second table entry in step 180. In this way, the table definition which was originally associated with two separate tables housing the first keyword is expanded by migrating the first keyword to a single integrated table definition of the present invention.

Accessing a unique row of a large database presents a number of performance issues, since often the table is indexed on its key, and if this key is modified regularly or comprises a large number of fields within the

table, performance may actually degrade such that access to a single row becomes a linear function where each row is checked to determine if a match has occurred. In these situations, the maximum number of key queries that the database would need to perform would be a function of the total number of rows within the table itself. For example, a query against a table having 10,000 rows would require at a maximum 10,000 checks against each row before a guaranteed result could be returned from the query.

As one skilled in the art would readily appreciate, database retrieval time should be as direct as possible and should not be a serial or linear function. Correspondingly, two tables may be created for purposes of the present invention, as depicted in Fig. 3, where the access to the first table 200 ("Table1") creates an index offset into a second table 300 ("Table2").

Table1 200 includes three fields: "ID" 210, "Keyword" 220, and "Value" 230. Values for the ID field 210 may be repeated within Table1 200, and need not be unique. However, values for the Keyword field 220 (e.g., First Keyword 260 through Nth Keyword 270) uniquely occur within Table1 200, as may values associated with the Value field 230 (e.g., First Value 280 through Zth Value 290). Although, since the Value field 230 is not part of the key, it is not required to be unique within Table 1 200. Accordingly, a key for Table1 200 may include only the ID 210 and the Keyword 220 field.

When a query to obtain a row of Table1 200 or Table2 300 is requested, the key is constructed using only the ID 210 and the Keyword 220 fields, with the smaller more compact Table1 200 being accessed initially and providing an index offset into the second larger Table2 300. This reduces the amount of search queries necessary to return a result,

since serial access begins, if at all, within Table2 300 and only where Keyword field values are not unique.

Table2 300 includes three fields: "ID" 310, "Keyword" 320, and "Value" 330. However, in Table2 300 the composite key to access a row of the table includes the ID 310, Keyword 320, and Value 330 fields. Correspondingly, Keyword values, such as First Keyword 350 may be repeated within Table2 and only a complete tuple of table 300 is ensured to be unique. As a result, the key to Table2 300 includes all fields of Table2 300 as presented.

As one skilled in the art will readily appreciate, by decomposing a table definition of the present invention into two separate tables with the first being smaller and more compact and being used to provide an offset access into a second much larger table, performance may be substantially enhanced. Especially, if values are continually being added to the second table because updates to the entire database would not need to occur as each change is made, and the smaller compact table would rarely need updated. Moreover, the number of serial accesses of each row of Table2 300 is limited to those instances in which Keyword values are repeated and not unique.

Furthermore by using two tables, as presented above, a database server software can more easily enforce integrity constraints to preclude multiple values for a key when such a situation is not permissible. In this way, the integrity of a key's cardinality may be checked by using this multiple table feature to enforce one-to-one or one-to-many cardinality relationships.

Fig. 4 depicts one flow diagram of an access method having two tables. A first and second table are created in step 400, with each table

having three fields I, K, and V. The first table is a smaller more compact table and initially the values associated with uniquely occurring I and K fields are inserted therein in step 410 where a first tuple of the first table having the values I1, K1, and V1. Concurrently, the first tuple is inserted
 5 into a second table in step 430, however uniquely occurring values in the second table for the I and K field are not required, rather the entire tuple is all that is guaranteed to be unique within the second table.

Subsequently, a second tuple I1, K1, and V2 are inserted into the second table in step 450. At some later point in time after the first and
 10 second table have been populated, a search request is received in step 420 for the tuple I1, K1, and V2. This request precipitates a search of table 1 in step 440 where tuple I1, K1, and V1 are detected using the key pair I1 and K1 to access the first table, this detection also provides an index offset into the second table in step 460, such that searching of the second table begins
 15 with the tuple I1, K1, and V1. The tuple I1, K1, and V2 are immediately detected in the second table and returned in step 470.

Fig. 5 depicts one flow diagram of a method for expanding a keyword. Initially, table T1 is received in step 480 having three fields associated therewith, namely I, K, and V. A key is created to access T1 in
 20 step 490, the key consisting on the tuple I, K, and V. Subsequently an insertion occurs into T1 where a K field has a value K1 repeated, and each individual K1 is associated with a different V value (e.g., V1 and V2, step 500). T1 may then be used to automatically construct a smaller table T2 which uses only uniquely occurring I and K values as its key.

25 Subsequently, a search may occur against table T2 in step 510 where a index offset is returned into T1 in step 520, and a further search

beginning at the offset location reveals a desired tuple value in step 530 which is returned in response to the initial search request.

The foregoing description of an exemplary embodiment of the invention has been presented for purposes of illustration and description. It

5 is not intended to be exhaustive nor to limit the invention to the precise form disclosed. Many alternatives, modifications, and variations will be apparent to those skilled in the art in light of the above teaching.

Accordingly, this invention is intended to embrace all alternatives, modifications, and variations that fall within the spirit and broad scope of
10 the attached claims.

11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65
66
67
68
69
70
71
72
73
74
75
76
77
78
79
80
81
82
83
84
85
86
87
88
89
90
91
92
93
94
95
96
97
98
99
100
101
102
103
104
105
106
107
108
109
110
111
112
113
114
115
116
117
118
119
120
121
122
123
124
125
126
127
128
129
130
131
132
133
134
135
136
137
138
139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158
159
160
161
162
163
164
165
166
167
168
169
170
171
172
173
174
175
176
177
178
179
180
181
182
183
184
185
186
187
188
189
190
191
192
193
194
195
196
197
198
199
200
201
202
203
204
205
206
207
208
209
210
211
212
213
214
215
216
217
218
219
220
221
222
223
224
225
226
227
228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
258
259
260
261
262
263
264
265
266
267
268
269
270
271
272
273
274
275
276
277
278
279
280
281
282
283
284
285
286
287
288
289
290
291
292
293
294
295
296
297
298
299
300
301
302
303
304
305
306
307
308
309
310
311
312
313
314
315
316
317
318
319
320
321
322
323
324
325
326
327
328
329
330
331
332
333
334
335
336
337
338
339
340
341
342
343
344
345
346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377
378
379
380
381
382
383
384
385
386
387
388
389
390
391
392
393
394
395
396
397
398
399
400
401
402
403
404
405
406
407
408
409
410
411
412
413
414
415
416
417
418
419
420
421
422
423
424
425
426
427
428
429
430
431
432
433
434
435
436
437
438
439
440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485
486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539
540
541
542
543
544
545
546
547
548
549
550
551
552
553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593
594
595
596
597
598
599
600
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647
648
649
650
651
652
653
654
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701
702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809
810
811
812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863
864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917
918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971
972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000